

Coded Cooperative Data Exchange for a Secret Key

Thomas A. Courtade

Dept. of Electrical Engineering and Computer Sciences
University of California, Berkeley
Email: courtade@eecs.berkeley.edu

Thomas R. Halford

TrellisWare Technologies
San Diego, CA
Email: thalford@trellisware.com

Abstract

We consider a coded cooperative data exchange problem with the goal of generating a secret key. Specifically, we investigate the number of public transmissions required for a set of clients to agree on a secret key with probability one, subject to the constraint that it remains private from an eavesdropper.

Although the problems are closely related, we prove that secret key generation with fewest number of linear transmissions is **NP**-hard, while it is known that the analogous problem in traditional cooperative data exchange can be solved in polynomial time. In doing this, we completely characterize the best-possible performance of linear coding schemes, and also prove that linear codes can be strictly suboptimal. Finally, we extend the single-key results to characterize the minimum number of public transmissions required to generate a desired integer number of statistically independent secret keys.

I. INTRODUCTION

In this paper, we consider a cooperative data exchange problem with the goal of generating a secret key. Primarily, we study the number of public transmissions required for a set of clients to agree on a secret key, subject to the constraint that it remains private from an eavesdropper.

In an asymptotic setting, the reciprocal relationship between secret key (SK) capacity and communication for omniscience was revealed in the pioneering work [2] by Csiszár and Narayan. They showed that the maximum rate at which secrecy can be generated by a collection of terminals is in one-to-one correspondence with the minimum rate required for those same terminals to communicate for omniscience. Though they characterized the minimum communication rate required to attain omniscience, Csiszár and Narayan left characterizing the minimum communication rate required to generate a maximum-rate SK as an open problem [2, Section VI].

In [3], El Rouayheb *et al.* introduced a non-asymptotic, combinatorial version of Csiszár and Narayan's communication for omniscience problem, which they called (*coded*) *cooperative data exchange* (CCDE). Since its introduction, this problem has received significant attention from many researchers (see [4] and the references therein). Algorithms and heuristics for solving the CCDE problem were presented in [5]–[7] for broadcast networks and in [8], [9] for multihop networks. Moreover, a number of authors

This work is supported by in part by the NSF Center for Science of Information under grant agreement CCF-0939370. This paper will appear in part at the 2014 International Symposium on Information Theory [1].

have considered generalizations of the CCDE problem to model various practical system considerations [10]–[14].

In [4], [8], Courtade and Wesel showed that the omniscience-secrecy relations in [2] translate nicely to the combinatorial CCDE setting. Building on [4], [8], we presently investigate the number of public transmissions required for a set of clients to agree on a SK with probability one, subject to the constraint that it remains private from an eavesdropper. In doing so, we address a combinatorial analog of Csiszár’s open problem. On this note, we remark that independent of the work on CCDE, Chan considered a closely related finite linear source model and gave suboptimal bounds on the public transmission block length required for perfect SK agreement [15]. Thus, our results give a definitive solution to Chan’s problem under the CCDE model.

Related to the present work is the minimum communication rate required to generate a maximum-rate SK in the asymptotic setting. In [16], Tyagi gave a multi-letter expression characterization for this rate in the two-terminal case in terms of the *r-rounds interactive common information*. Recently, Mukherjee and Kshyap considered extensions to the multi-terminal case [17].

Despite the similarity in spirit, the asymptotic setting of [16] and the combinatorial setting of the present paper are considerably different in nature, and the proof techniques used are orthogonal. That said, all of these results shed light on the fundamentally different natures of SK generation at minimum communication rate and SK generation via communication for omniscience.

The weakly secure CCDE problem introduced in [18] is also related to our work. The goal of the weakly secure CCDE problem is to communicate for omniscience while revealing as little information as possible to an eavesdropper. This is closely related to the CCDE under a privacy constraint problem studied in [4]. Yan and Sprintson designed coding schemes that solve the weakly secure CCDE problem while revealing as little as information as possible to an eavesdropper [18]. Improvements to these schemes can be derived using the codes described in [19]. The primary distinction between the present setting and that of the weakly secure CCDE problem is that we only aim to generate a SK; we do not require that the nodes communicate for omniscience nor do we require that the SK corresponds to any given message.

Finally, we remark that in a recent paper [20], Halford *et al.* developed practical protocols for SK generation in ad hoc networks based on the CCDE problem. Briefly, a scenario was studied wherein the protocol designer controls the initial distribution of master keys so that secret keys can later be efficiently generated among arbitrary groups of clients. The results given in the present paper establish limits and suggest design rules for such protocols.

Our Contributions

Given the close connection between CCDE and SK generation, we show two surprising results. First, we prove that finding an optimal (linear) coding scheme for SK generation is NP-hard, while it is known that the analogous CCDE problem is in P. In doing this, we completely characterize the attainable performance for linear coding schemes in terms of hypergraph connectivity. Second, despite linear codes being optimal for CCDE, we demonstrate that they can be strictly suboptimal for SK generation. Several ancillary results are also proved.

This paper is organized as follows. Section II formally defines our system model and reviews relevant results on CCDE. In Section III, we state and prove our main results for the generation of a single SK. Section IV characterizes the minimum number of public transmissions required to generate multiple SKs, and Section V delivers concluding remarks.

II. SYSTEM MODEL AND PRELIMINARIES

We first establish basic notation. Throughout, we use calligraphic notation to denote sets. For two sets $\mathcal{A} \subset \mathcal{B}$, we write $\mathcal{B} \setminus \mathcal{A}$ to denote those elements in \mathcal{B} , but not in \mathcal{A} . If \mathcal{A} is a singleton set (i.e., $\mathcal{A} = \{a\}$), then we often use the notation $\mathcal{B} - a \triangleq \mathcal{B} \setminus \{a\}$ for convenience. We define \mathbb{Z} to be the set of integers. For positive $m \in \mathbb{Z}$, we use the shorthand notation $[m] \triangleq \{1, 2, \dots, m\}$. Finally, for random variables X, Y , we write $I(X; Y)$ for the mutual information between X and Y .

A. System Model

Throughout, we consider networks defined by a set of n clients (i.e., terminals) $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$, a positive integer m , and a family of finite sets $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$ (each $\mathcal{I}_j \subseteq [m]$ and $\cup_{j=1}^n \mathcal{I}_j = [m]$) in the following way. Define the random (column) vector $\underline{X} \triangleq [X_1, X_2, \dots, X_m]^T$, where each X_i is a discrete random variable with equiprobable distribution on a finite field \mathbb{F} , and (X_1, X_2, \dots, X_m) are mutually independent¹. The random variables $\{X_i\}_{i=1}^m$ are called *messages*, and $\{X_i : i \in \mathcal{I}_j\}$ is the set of messages initially held by client $c_j \in \mathcal{C}$. In other words, \mathcal{I}_j defines the indices of messages initially held by client c_j , for $j = 1, \dots, n$. Throughout, n will always denote the number of clients; since the sets \mathcal{I}_j are always indexed by $j \in [n]$, we will use the shorthand notation $\{\mathcal{I}_j\}$ to denote the family $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$.

We adopt the communication model which is standard in index coding and CCDE problems. That is, we consider transmission schemes consisting of a finite number of communication rounds. In each round, a single client broadcasts an element of \mathbb{F} (which can be a function of the messages initially held by that client and all previous transmissions) to all other clients over an error-free channel. It is further assumed that all clients have knowledge of the index sets $\mathcal{I}_1, \dots, \mathcal{I}_n$, and thus follow a protocol which is mutually agreed upon. We will elaborate on the definition of a transmission protocol in the next subsection.

B. Transmission Protocols Defined

For a network defined by $\{\mathcal{I}_j\}$, a transmission protocol \mathbf{P} (or simply, a protocol \mathbf{P}) consisting of t communication rounds is defined by n encoding functions $\{f_1, f_2, \dots, f_n\}$, and a t -tuple (i_1, i_2, \dots, i_t) , where $i_k \in [n]$ indicates which client transmits during communication round k . More specifically, during communication round k , client c_{i_k} transmits

$$f_{i_k}(\{X_j : j \in \mathcal{I}_{i_k}\}, k, \{f_{i_\ell}\}_{\ell=1}^{k-1}) \in \mathbb{F}, \quad (1)$$

where we have abbreviated the transmitted symbols in rounds $\ell \in [k-1]$ by $\{f_{i_\ell}\}_{\ell=1}^{k-1}$. For a given transmission protocol \mathbf{P} requiring t communication rounds, we let $\mathbf{T}(\underline{X}, \mathbf{P}) \in \mathbb{F}^t$ be the column vector with k^{th} entry equal to $f_{i_k}(\{X_j : j \in \mathcal{I}_{i_k}\}, k, \{f_{i_\ell}\}_{\ell=1}^{k-1})$. Letting $\|\cdot\|$ be the length function, we have

¹For technical reasons, we assume $|\mathbb{F}| > n$.

$\|\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})\| = t$. Note that $\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})$ is a random variable since it is a function of the random vector $\underline{\mathbf{X}}$. Generally, the transmission protocol under consideration will be clear from context. Hence, we abbreviate $\mathbf{T}(\underline{\mathbf{X}}) \triangleq \mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})$ for convenience when there is no ambiguity.

A transmission protocol is said to be *linear* (over \mathbb{F}) if the encoding functions $\{f_1, f_2, \dots, f_n\}$ are of the form

$$f_{i_k}(\{X_j : j \in \mathcal{I}_{i_k}\}, k, \{f_{i_\ell}\}_{\ell=1}^{k-1}) = \sum_j \alpha_j^{(k)} X_j, \quad (2)$$

where $\alpha_j^{(k)} \in \mathbb{F}$ can be interpreted as the encoding coefficient for message j during communication round k . In this case, we can express $\mathbf{T}(\underline{\mathbf{X}}) = A\underline{\mathbf{X}}$, where $A \in \mathbb{F}^{t \times m}$ assuming the definitions $t \triangleq \|\mathbf{T}(\underline{\mathbf{X}})\|$ and $m \triangleq |\cup_j \mathcal{I}_j|$. Hence, the *encoding matrix* A provides a succinct description of a linear transmission protocol. Note that the order of transmissions corresponding to a linear protocol is inconsequential.

C. Transmission Protocols for Omniscience

A transmission protocol \mathbf{P} is said to achieve *omniscience* if there exist decoding functions $\{g_1, g_2, \dots, g_n\}$ which satisfy

$$g_j(\{X_i : i \in \mathcal{I}_j\}, \mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})) = \underline{\mathbf{X}} \text{ for each } j \in [n] \quad (3)$$

with probability 1.

Before proceeding, let $M^*(\{\mathcal{I}_j\})$ denote the optimal value of the following integer linear program (ILP):

$$\begin{aligned} & \text{minimize : } \sum_{j \in [n]} a_j \\ & \text{subject to : } \sum_{j \in \mathcal{S}} a_j \geq \left| \bigcap_{j \in \bar{\mathcal{S}}} \bar{\mathcal{I}}_j \right| \text{ for all nonempty } \mathcal{S} \subset [n] \\ & \quad a_j \in \mathbb{Z} \text{ for all } j \in [n], \end{aligned} \quad (4)$$

where $\bar{\mathcal{I}}_i \triangleq (\cup_j \mathcal{I}_j) \setminus \mathcal{I}_i$ and $\bar{\mathcal{S}} \triangleq [n] \setminus \mathcal{S}$. The quantity $M^*(\{\mathcal{I}_j\})$ will play an important role in our treatment due to its inherent connection to the communication for omniscience, which is made explicit by the following theorem².

Theorem 1. [4, Theorem 2] *If a protocol \mathbf{P} achieves omniscience, then $\|\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})\| \geq M^*(\{\mathcal{I}_j\})$. Conversely, there always exists a linear protocol \mathbf{P}_L that achieves omniscience and has $\|\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P}_L)\| = M^*(\{\mathcal{I}_j\})$.*

Theorem 1 addresses the central issue in the CCDE problem, which primarily investigates the number of transmissions required to achieve omniscience. We remark that this is not equivalent to characterizing the minimum communication *rate* required for omniscience (as would be the case in the original communication for omniscience problem [2]) due in part to the integrality constraint on the number of transmissions.

²Theorem 1 essentially appeared in the given form in [8]. However, it was independently discovered by Milosavljevic et al. [11] and Chan [15] at roughly the same time.

D. Transmission Protocols for Secret Keys

A transmission protocol (with corresponding transmission sequence $\mathbf{T}(\underline{\mathbf{X}})$) generates a secret key (SK) if there exist decoding functions $\{k_1, k_2, \dots, k_n\}$ which satisfy the following three properties:

- (i) For all $j \in [n]$, and with probability 1,

$$k_j(\{X_i : i \in \mathcal{I}_j\}, \mathbf{T}(\underline{\mathbf{X}})) = k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}})).$$

- (ii) $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$ is equiprobable on \mathbb{F} .

- (iii) $I(k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}})); \mathbf{T}(\underline{\mathbf{X}})) = 0$.

In words, requirement (iii) guarantees that the public transmissions $\mathbf{T}(\underline{\mathbf{X}})$ reveal no information about $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$. Requirement (i) asserts that all clients $c_j \in \mathcal{C}$ can compute $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$. For these reasons, $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$ is called a *secret key*. Naturally, a secret key should be equiprobable on its domain to make guessing difficult, thus motivating requirement (ii).

It is not immediately clear whether any protocol \mathbf{P} generates a SK. However, it turns out that such protocols exist in great abundance. In particular, the existence of protocols that generate a SK depends solely on the family $\{\mathcal{I}_j\}$.

Theorem 2. [4, Theorem 6] *For a network defined by $\{\mathcal{I}_j\}$, there exists a protocol \mathbf{P} which generates a SK if and only if*

$$|\cup_j \mathcal{I}_j| \geq M^*(\{\mathcal{I}_j\}) + 1. \quad (5)$$

Despite the fact that $M^*(\{\mathcal{I}_j\})$ corresponds to the optimal value of an ILP, it can be computed in time polynomial in the number of messages $m = |\cup_j \mathcal{I}_j|$ (see [4], [11]). Therefore, for any family $\{\mathcal{I}_j\}$, we can efficiently test whether (5) holds. Hence, the essential remaining question is: “How many transmissions are needed to generate a SK?”

To this end, let $\mathcal{P}(\{\mathcal{I}_j\})$ denote the set of protocols for $\{\mathcal{I}_j\}$ that generate a SK, and define

$$S(\{\mathcal{I}_j\}) \triangleq \min \left\{ \|\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})\| : \mathbf{P} \in \mathcal{P}(\{\mathcal{I}_j\}) \right\}. \quad (6)$$

That is, $S(\{\mathcal{I}_j\})$ is the minimum number of transmissions needed to generate a SK. Similarly, let $\mathcal{P}_L(\{\mathcal{I}_j\})$ denote the set of linear protocols for $\{\mathcal{I}_j\}$ that generate a SK, and define

$$S_L(\{\mathcal{I}_j\}) \triangleq \min \left\{ \|\mathbf{T}(\underline{\mathbf{X}}, \mathbf{P})\| : \mathbf{P} \in \mathcal{P}_L(\{\mathcal{I}_j\}) \right\}. \quad (7)$$

In words, $S_L(\{\mathcal{I}_j\})$ is the minimum number of transmissions required to generate a SK when we restrict our attention to linear protocols. If $\{\mathcal{I}_j\}$ does not satisfy (5), then we set $S(\{\mathcal{I}_j\}) = S_L(\{\mathcal{I}_j\}) = \infty$.

Remark 1. We will often write “ $\{\mathcal{I}_j\}$ generates a SK” instead of the more accurate, but cumbersome, “For the network defined by $\{\mathcal{I}_j\}$, there exists a protocol \mathbf{P} which generates a SK” whenever (5) holds.

III. GENERATING A SINGLE SECRET KEY

In this section, we investigate the number of transmissions required to generate a SK. In particular, we completely characterize $S_L(\{\mathcal{I}_j\})$, and make progress toward characterizing $S(\{\mathcal{I}_j\})$. We will treat the more general case of generating multiple secret keys with minimum public communication in Section IV. Since the single-SK setting is arguably the most important in practice and the notation is less cumbersome than the general case, we find it beneficial to highlight the single-SK setting in the present section.

As demonstrated in the previous section, the CCDE and SK-generation problems are closely connected through the quantity $M^*(\{\mathcal{I}_j\})$. Since Theorem 1 and the tractability of ILP (4) essentially resolve the CCDE problem, it is natural to conjecture that a similar result should hold for $S(\{\mathcal{I}_j\})$ and $S_L(\{\mathcal{I}_j\})$. Unfortunately, there is a fundamental difference between the problems, which is revealed by the following two negative results:

Theorem 3. *Computing $S_L(\{\mathcal{I}_j\})$ is NP-hard.*

Theorem 4. *For any integer k , there exist families $\{\mathcal{I}_j\}$ for which $S_L(\{\mathcal{I}_j\}) > S(\{\mathcal{I}_j\}) + k$.*

For the CCDE problem, Theorem 1 asserts that linear protocols achieve optimal performance. Furthermore, the number of transmissions required by linear protocols is easily computed. For the problem of SK generation, the opposite is true. That is, linear protocols can be suboptimal, and the number of transmissions required by linear protocols is generally difficult to compute. This situation is parallel to that of multicast network coding and index coding. The two problems are closely related (cf. [21]), but exhibit the same dichotomy. See [22]–[24] and our remark at the end of this section for more details.

A. Proof of Theorem 3

Despite the negative results offered by Theorems 3 and 4, we can characterize several properties of $S_L(\{\mathcal{I}_j\})$, $S(\{\mathcal{I}_j\})$, and $M^*(\{\mathcal{I}_j\})$. Some of these properties are demonstrated in the following results, which are needed as we progress toward proving Theorem 3. A complete characterization of $S_L(\{\mathcal{I}_j\})$ will be given in Theorem 5.

Lemma 1. *If $\{\mathcal{I}_j\}$ generates a SK, then*

$$S(\{\mathcal{I}_j\}) \leq S_L(\{\mathcal{I}_j\}) \leq M^*(\{\mathcal{I}_j\}). \quad (8)$$

Proof: By definition, $S(\{\mathcal{I}_j\}) \leq S_L(\{\mathcal{I}_j\})$ since $\mathcal{P}_L(\{\mathcal{I}_j\}) \subseteq \mathcal{P}(\{\mathcal{I}_j\})$. The second inequality follows from the proof of [4, Theorem 6], in which a linear transmission protocol \mathbf{P} is constructed that generates a SK with $\|\mathbf{T}(\mathbf{X}), \mathbf{P}\| = M^*(\{\mathcal{I}_j\})$ communication rounds. ■

We say that $\{\mathcal{J}_j\}$ is a *subfamily* of $\{\mathcal{I}_j\}$ if there is a set $\mathcal{S} \subset \cup_j \mathcal{I}_j$ such that $\mathcal{J}_j = \mathcal{I}_j \setminus \mathcal{S}$ for all $j \in [n]$.

Lemma 2. *If $\{\mathcal{J}_j\}$ is a subfamily of $\{\mathcal{I}_j\}$, then*

$$M^*(\{\mathcal{J}_j\}) \leq M^*(\{\mathcal{I}_j\}), \quad (9)$$

$$S(\{\mathcal{J}_j\}) \geq S(\{\mathcal{I}_j\}), \text{ and} \quad (10)$$

$$S_L(\{\mathcal{J}_j\}) \geq S_L(\{\mathcal{I}_j\}). \quad (11)$$

Proof: By De Morgan's law, it is easy to verify that

$$\left| \bigcap_{j \in \mathcal{S}} \bar{\mathcal{I}}_j \right| \geq \left| \bigcap_{j \in \mathcal{S}} \bar{\mathcal{J}}_j \right| \quad \text{for all nonempty } \mathcal{S} \subset [n], \quad (12)$$

where $\bar{\mathcal{I}}_i \triangleq (\cup_j \mathcal{I}_j) \setminus \mathcal{I}_i$ and $\bar{\mathcal{J}}_i \triangleq (\cup_j \mathcal{J}_j) \setminus \mathcal{J}_i$. Therefore, the constraints in ILP (4) are relaxed, and $M^*(\{\mathcal{J}_j\}) \leq M^*(\{\mathcal{I}_j\})$ by definition.

To show (10), observe that any transmission protocol which generates a SK for the subfamily $\{\mathcal{J}_j\}$ also generates a SK for the family $\{\mathcal{I}_j\}$ by ignoring the set of messages $\{X_i : i \notin \cup_j \mathcal{J}_j\}$. Hence, it follows that $S(\{\mathcal{J}_j\}) \geq S(\{\mathcal{I}_j\})$. If $\{\mathcal{J}_j\}$ can not generate a SK, the inequality trivially holds. This argument also proves (11). ■

Lemma 2 demonstrates monotonicity, but offers no insight into whether inequalities (9)-(11) are tight. The following lemma identifies settings under which (11) holds with equality, and will prove useful later on.

Lemma 3. *If $S_L(\{\mathcal{I}_j\}) < M^*(\{\mathcal{I}_j\})$, then there exists some $\ell \in \cup_j \mathcal{I}_j$ for which $S_L(\{\mathcal{I}_j - \ell\}) = S_L(\{\mathcal{I}_j\})$.*

Proof: Define $m \triangleq |\cup_j \mathcal{I}_j|$. By definition, there is a linear transmission protocol \mathbf{P}_L which generates a SK in $S_L(\{\mathcal{I}_j\})$ communication rounds. Let $\mathbf{T}(\mathbf{X}, \mathbf{P}_L) = A\mathbf{X}$ be the sequence of transmissions made by \mathbf{P}_L , and let $\{k_1, \dots, k_n\}$ be valid decoding functions.

Since $\|\mathbf{T}(\mathbf{X}, \mathbf{P}_L)\| = S_L(\{\mathcal{I}_j\}) < M^*(\{\mathcal{I}_j\})$, Theorem 1 asserts that the protocol \mathbf{P}_L can not achieve omniscience. Therefore, by a possible permutation of clients, we can assume without loss of generality that there is no function g_1 for which

$$g_1(\{X_i : i \in \mathcal{I}_1\}, A\mathbf{X}) = \mathbf{X} \quad \text{with probability 1.} \quad (13)$$

As a consequence, there must exist a nonzero vector \mathbf{v} such that $A\mathbf{v} = 0$, and $v_i = 0$ for all $i \in \mathcal{I}_1$. Indeed, if there is no such \mathbf{v} , then A has empty nullspace and client c_1 can solve a full-rank system of equations to recover \mathbf{X} , yielding a contradiction. Since \mathbf{v} is not identically zero, there is some $\ell \notin \mathcal{I}_1$ for which $v_\ell \neq 0$.

Considering any such ℓ , we define $\hat{X}_\ell \equiv 0$, and $\hat{X}_i \triangleq X_i$ for $i \in \cup_j (\mathcal{I}_j - \ell)$. Also, define vectors $\hat{\mathbf{X}} \triangleq [\hat{X}_1, \hat{X}_2, \dots, \hat{X}_m]^T$ and $\mathbf{X}' \triangleq \hat{\mathbf{X}} + X_\ell \cdot \mathbf{v}$. First, we note that

$$k_j(\{\hat{X}_i : i \in \mathcal{I}_j\}, A\hat{\mathbf{X}}) = k_1(\{\hat{X}_i : i \in \mathcal{I}_1\}, A\hat{\mathbf{X}})$$

for all $j \in [n]$ since

$$k_j(\{X_i : i \in \mathcal{I}_j\}, A\mathbf{X}) = k_1(\{X_i : i \in \mathcal{I}_1\}, A\mathbf{X})$$

with probability 1 by definition.

Next, observe that:

$$I(k_1(\{\hat{X}_i : i \in \mathcal{I}_1\}, A\hat{\mathbf{X}}); A\hat{\mathbf{X}}) = I(k_1(\{\hat{X}_i + X_\ell \cdot v_i : i \in \mathcal{I}_1\}, A\mathbf{X}'); A\mathbf{X}') \quad (14)$$

$$= I(k_1(\{X_i : i \in \mathcal{I}_1\}, A\mathbf{X}); A\mathbf{X}) = 0 \quad (15)$$

In the above,

- (14) follows since $A\hat{\underline{X}}' = A(\hat{\underline{X}} + X_\ell \cdot \underline{v}) = A\hat{\underline{X}}$, and $v_i = 0$ for all $i \in \mathcal{I}_1$.
- (15) follows from the (crucial) observation that $\hat{\underline{X}}'$ and $\hat{\underline{X}}$ are equal in distribution, and by definition of A and k_1 .

Finally, by similar reasoning, we note that the random variable $k_1(\{\hat{X}_i : i \in \mathcal{I}_1\}, A\hat{\underline{X}})$ is equiprobable on \mathbb{F} since

$$k_1(\{\hat{X}_i : i \in \mathcal{I}_1\}, A\hat{\underline{X}}) = k_1(\{\hat{X}_i + X_\ell \cdot v_i : i \in \mathcal{I}_1\}, A\hat{\underline{X}}') \quad (16)$$

$$\stackrel{d}{=} k_1(\{X_i : i \in \mathcal{I}_1\}, A\hat{\underline{X}}), \quad (17)$$

and $k_1(\{X_i : i \in \mathcal{I}_1\}, A\hat{\underline{X}})$ is equiprobable on \mathbb{F} by definition. In (17), the notation $\stackrel{d}{=}$ indicates equality in distribution.

Therefore, we can conclude that a SK can be generated by the subfamily $\{\mathcal{I}_j - \ell\}_j$ by applying the protocol \mathbf{P}_L and fixing $X_\ell \equiv 0$. This proves that $S_L(\{\mathcal{I}_j - \ell\}) \leq S_L(\{\mathcal{I}_j\})$. By Lemma 2, the reverse inequality also holds. ■

In order to proceed, we will need to introduce *critical* families. To this end, let $\tau \geq 1$ be an integer. A family $\{\mathcal{I}_j\}$ is τ -critical if the following hold:

- (i) $|\cup_j \mathcal{I}_j| - M^*(\{\mathcal{I}_j\}) = \tau$, and
- (ii) $M^*(\{\mathcal{I}_j - i\}) = M^*(\{\mathcal{I}_j\})$ for all $i \in \cup_j \mathcal{I}_j$.

It is interesting to note that τ -criticality of $\{\mathcal{I}_j\}$ can be efficiently tested since $M^*(\{\mathcal{I}_j\})$ is computable in polynomial time. Observe that 1-critical families enjoy a threshold property: families $\{\mathcal{I}_j\}$ that are 1-critical generate a secret key, and no proper subfamilies of $\{\mathcal{I}_j\}$ generate SKs. This is a consequence of Theorem 2 and the definition of 1-criticality.

A *minimum τ -critical subfamily* $\{\mathcal{J}_j^*\}$ of $\{\mathcal{I}_j\}$ satisfies

$$|\cup_j \mathcal{J}_j^*| \leq |\cup_j \mathcal{I}_j| \quad (18)$$

for all other τ -critical subfamilies $\{\mathcal{J}_j\}$ of $\{\mathcal{I}_j\}$. Note that if $\{\mathcal{I}_j\}$ is 1-critical, then $\{\mathcal{I}_j\}$ is its own unique minimum 1-critical subfamily.

The following Theorem demonstrates that minimum 1-critical subfamilies completely characterize $S_L(\{\mathcal{I}_j\})$.

Theorem 5. *If $\{\mathcal{I}_j\}$ generates a SK, then*

$$S_L(\{\mathcal{I}_j\}) = M^*(\{\mathcal{J}_j^*\}) = |\cup_j \mathcal{J}_j^*| - 1, \quad (19)$$

where $\{\mathcal{J}_j^*\}$ is a minimum 1-critical subfamily of $\{\mathcal{I}_j\}$.

Proof: By inductively applying Lemma 3, we can find a subfamily $\{\mathcal{T}_j\}$ of $\{\mathcal{I}_j\}$ for which

$$S_L(\{\mathcal{I}_j\}) = M^*(\{\mathcal{T}_j\}). \quad (20)$$

Let $\{\mathcal{I}_j\}$ be any 1-critical subfamily of $\{\mathcal{T}_j\}$. We have the following chain of inequalities

$$S_L(\{\mathcal{I}_j\}) \leq S_L(\{\mathcal{I}_j^*\}) \leq M^*(\{\mathcal{I}_j^*\}) \quad (21)$$

$$\leq M^*(\{\mathcal{I}_j\}) \quad (22)$$

$$\leq M^*(\{\mathcal{T}_j\}) \quad (23)$$

$$= S_L(\{\mathcal{I}_j\}). \quad (24)$$

The above steps can be justified as follows:

- (21) follows from Lemmas 1 and 2.
- By definition of τ -criticality, (18) is equivalent to $M^*(\{\mathcal{I}_j^*\}) \leq M^*(\{\mathcal{I}_j\})$. Thus, (22) follows since $\{\mathcal{I}_j^*\}$ is a minimum 1-critical subfamily of $\{\mathcal{I}_j\}$, and $\{\mathcal{I}_j\}$ is a 1-critical subfamily of $\{\mathcal{I}_j\}$.
- (23) follows from Lemma 2.
- (24) is the assertion of (20).

This proves that $S_L(\{\mathcal{I}_j\}) = M^*(\{\mathcal{I}_j^*\})$. Recalling the definition of 1-criticality completes the proof. ■

The network defined by $\{\mathcal{I}_j\}$ has a natural representation as a hypergraph³. In particular, we make the following definition:

Definition 1. Consider a hypergraph $H = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V} = \mathcal{C}$, and edge set $\mathcal{E} = \cup_j \mathcal{I}_j$. H is the hypergraph representation of $\{\mathcal{I}_j\}$ iff it has the following property: a vertex $c_j \in \mathcal{V}$ is contained in the edge $e \in \mathcal{E}$ if and only if $e \in \mathcal{I}_j$.

Theorem 5 implies that $S_L(\{\mathcal{I}_j\})$ is easily computed if we can identify a minimum 1-critical subfamily of $\{\mathcal{I}_j\}$. By Theorem 3, we know this must be NP-hard. In order to prove this to be the case, we require the following lemma which lends a hypergraph interpretation to 1-criticality. For a hypergraph $H = (\mathcal{V}, \mathcal{E})$, an edge set $\mathcal{E}' \subseteq \mathcal{E}$ is a *minimal connected dominating edge set* if the subhypergraph $H' = (\mathcal{V}, \mathcal{E}')$ is connected, and the removal of any edge from \mathcal{E}' disconnects H' .

Lemma 4. Let $H = (\mathcal{V}, \mathcal{E})$ be the hypergraph representation of $\{\mathcal{I}_j\}$. H is connected if and only if

$$M^*(\{\mathcal{I}_j\}) < |\cup_j \mathcal{I}_j|. \quad (25)$$

In particular, $\{\mathcal{I}_j\}$ is 1-critical if and only if \mathcal{E} is a minimal connected dominating edge set.

Proof: First, suppose H is not connected. By definition, there must exist a nontrivial partition $\mathcal{V} = (\mathcal{S}, \bar{\mathcal{S}})$ such that there is no edge $e \in \mathcal{E}$ which contains vertices from both \mathcal{S} and $\bar{\mathcal{S}}$. Stated another way, $(\cup_{j \in \mathcal{S}} \mathcal{I}_j) \cap (\cup_{j \in \bar{\mathcal{S}}} \mathcal{I}_j) = \emptyset$. Hence, ILP (4) includes the two constraints

$$\sum_{j \in \mathcal{S}} a_j \geq \left| \bigcap_{j \in \bar{\mathcal{S}}} \bar{\mathcal{I}}_j \right| = \left| \bigcup_{j \in \mathcal{S}} \mathcal{I}_j \right| \quad (26)$$

$$\sum_{j \in \bar{\mathcal{S}}} a_j \geq \left| \bigcap_{j \in \mathcal{S}} \bar{\mathcal{I}}_j \right| = \left| \bigcup_{j \in \bar{\mathcal{S}}} \mathcal{I}_j \right|, \quad (27)$$

³We adopt the definition of a hypergraph that allows for repeated edges (i.e., multiple edges, with the same set of vertices, are permitted).

the sum of which imply $M^*(\{\mathcal{I}_j\}) \geq |\cup_j \mathcal{I}_j|$. By taking the contrapositive, we have proven

$$M^*(\{\mathcal{I}_j\}) < |\cup_j \mathcal{I}_j| \implies H \text{ is connected.} \quad (28)$$

Next, suppose H is connected, and assume without loss of generality that $\mathcal{E} = \cup_j \mathcal{I}_j \triangleq \{1, 2, \dots, m\}$. Since H is connected, there is a transmission protocol for which the entries of $\mathbf{T}(\underline{X})$ are precisely $\{X_1 + X_j\}_{j=2}^m$. Indeed, by connectivity of H , there must be some client c initially holding X_1 and some X_e (say, X_2 without loss of generality), and can therefore transmit $X_1 + X_2$ during the first communication round. By induction, assume that $\{X_1 + X_j\}_{j=2}^{m-1}$ are transmitted during the first $m-2$ communication rounds (permuting indices of the X_i 's if necessary). Again, by connectivity of H , there must be a client c' which initially holds X_m and X_k , where $k < m$. Hence, in communication round $m-1$, client c' can transmit $(X_1 + X_k) - (X_k - X_m) = X_1 + X_m$. Noting that

$$(X_1, X_1 + X_2, \dots, X_1 + X_m) \stackrel{d}{=} (X_1, X_2, \dots, X_m),$$

we have $I(X_1; \mathbf{T}(\underline{X})) = 0$. If client $c \in e \in \mathcal{E}$, then it can recover X_1 from the transmission $X_1 + X_e$ by simply subtracting X_e . Since H is connected, each $c \in \mathcal{V}$ belongs to some edge in \mathcal{E} , and therefore all clients can recover X_1 losslessly. Since X_1 is equiprobable on \mathbb{F} by definition, we can conclude that $\{\mathcal{I}_j\}$ generates a SK. Theorem 2 asserts that we must have $M^*(\{\mathcal{I}_j\}) < |\cup_j \mathcal{I}_j|$, and we have proven

$$M^*(\{\mathcal{I}_j\}) < |\cup_j \mathcal{I}_j| \iff H \text{ is connected.} \quad (29)$$

We now prove the second claim. To this end, suppose $\{\mathcal{I}_j\}$ is 1-critical. Then $M^*(\{\mathcal{I}_j\}) = |\cup_j \mathcal{I}_j| - 1$, which implies H is connected (and thus \mathcal{E} is dominating) by (55). Consider the subhypergraph $H' = (\mathcal{V}, \mathcal{E} \setminus \{e\})$, which corresponds to the subfamily $\{\mathcal{I}_j - e\}$ of $\{\mathcal{I}_j\}$. Since $\{\mathcal{I}_j\}$ is 1-critical, we must have $M^*(\{\mathcal{I}_j - e\}) = M^*(\{\mathcal{I}_j\}) = |\cup_j \mathcal{I}_j| - 1 = |\cup_j (\mathcal{I}_j - e)|$. By (55), H' must be disconnected, and therefore \mathcal{E} is a minimal connected dominating edge set.

On the other hand, suppose \mathcal{E} is a minimal connected dominating edge set. Since H is connected, (55) implies

$$M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - 1. \quad (30)$$

Since \mathcal{E} is minimal, for any $e \in \mathcal{E}$, $H' = (\mathcal{V}, \mathcal{E} \setminus \{e\})$ is disconnected, and (55) implies

$$M^*(\{\mathcal{I}_j - e\}) \geq |\cup_j (\mathcal{I}_j - e)| = |\cup_j \mathcal{I}_j| - 1. \quad (31)$$

Applying Lemma 2, we must have $M^*(\{\mathcal{I}_j\}) = M^*(\{\mathcal{I}_j - e\})$, and $|\cup_j \mathcal{I}_j| - M^*(\{\mathcal{I}_j\}) = 1$, which implies $\{\mathcal{I}_j\}$ is 1-critical. ■

We are finally in a position to prove Theorem 3.

Proof of Theorem 3: Let $H = (\mathcal{V}, \mathcal{E})$ be the hypergraph representation of $\{\mathcal{I}_j\}$. We can assume $\{\mathcal{I}_j\}$ generates a SK. By Theorem 5 and Lemma 4, computing $S_L(\{\mathcal{I}_j\})$ is equivalent to computing the number of edges in a minimum connected dominating edge set (i.e., a minimal connected dominating edge set with fewest possible edges). It is easy to see that the NP-complete SET COVER DECISION PROBLEM is a special case.

Indeed, consider any subsets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ whose union covers a finite set \mathcal{U} . For $u' \notin \mathcal{U}$, define $\mathcal{U}' = \mathcal{U} \cup \{u'\}$, and $\mathcal{A}'_j = \mathcal{A}_j \cup \{u'\}$ for $j \in [k]$. Clearly, $\{\mathcal{A}_{j_i}\}_{i=1}^m$ is a minimum cover of \mathcal{U} if and only if $\{\mathcal{A}'_{j_i}\}_{i=1}^m$ is a minimum connected cover of \mathcal{U}' . ■

Remark 2. Together, Theorem 5 and Lemma 4 give a succinct characterization of $S_L(\{\mathcal{I}_j\})$ in terms of hypergraph connectivity. We extend this result to the generation of multiple secret keys at the end of Section IV using a stronger form of hypergraph connectivity.

B. Proof of Theorem 4

Before proving Theorem 4, consider the following constructive example: Let $n = 7$, and consider the family $\{\mathcal{I}_j\}$ defined by $\mathcal{I}_1 = \{1, 2, 3, 4\}$, and $\mathcal{I}_2, \dots, \mathcal{I}_7$ are all $\binom{4}{2}$ distinct 2-element subsets of $\{1, 2, 3, 4\}$. By direct computation, we find that $\{\mathcal{I}_j - \{1\}\}$ is a minimum 1-critical subfamily, and hence $S_L(\{\mathcal{I}_j\}) = 2$ by Theorem 5. Suppose $\mathbb{F} = \{0, 1, \alpha, \beta\}^2 = \text{GF}(4) \times \text{GF}(4)$. Thus, we can express $X_j = (X_j^{(1)}, X_j^{(2)})$ for each $j = 1, \dots, 4$, where $X_j^{(1)}, X_j^{(2)}$ are mutually independent, each equiprobable on $\text{GF}(4)$. It is readily verified that the single transmission

$$\left(X_1^{(1)} + \alpha X_2^{(1)} + X_3^{(1)}, X_1^{(1)} + \beta X_2^{(1)} + X_4^{(1)} \right) \in \mathbb{F} \quad (32)$$

by client c_1 permits reconstruction of the SK

$$k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\mathbf{X})) = (X_3^{(1)}, X_4^{(1)}) \in \mathbb{F} \quad (33)$$

at all clients. Hence, we can conclude $1 = S(\{\mathcal{I}_j\}) < S_L(\{\mathcal{I}_j\}) = M^*(\{\mathcal{I}_j\}) = 2$.

The above construction is a *vector-linear* transmission protocol, and cannot be realized by a protocol which is linear over \mathbb{F} . A natural question is whether it is possible to bound the gap between $S(\{\mathcal{I}_j\})$ and $S_L(\{\mathcal{I}_j\})$. As asserted by Theorem 4, the answer to this is negative. Indeed, it is straightforward to generalize the previous construction and make the gap arbitrarily large.

To this end, consider a network of $n = \binom{m}{2} + 1$ clients such that $\mathcal{I}_1 = [m]$ and the other $\binom{m}{2}$ clients possess distinct pairs of messages. Observe that the 1-critical subfamilies of $\{\mathcal{I}_j\}$ are obtained by removing a single message – i.e., if any two messages $m_1, m_2 \in [m]$ are removed then the resulting hypergraph representation of $\{\mathcal{I}_j - \{m_1, m_2\}\}$ is no longer connected. This implies that $S_L(\{\mathcal{I}_j\}) = m - 2$. To show that there exists a nonlinear scheme that can do better, we show that $M^*(\{\mathcal{I}_j\}) = m - 2$:

- To show that $M^*(\{\mathcal{I}_j\}) \leq m - 2$, let client c_1 transmit $m - 2$ independent linear combinations of the messages. Provided the encoding matrix A is full rank (e.g., a Vandermonde matrix), every other node can use its own pair of messages to recover the other $m - 2$.
- We note that $M^*(\{\mathcal{I}_j\}) \geq S_L(\{\mathcal{I}_j\}) = m - 2$ by Lemma 1, and therefore $M^*(\{\mathcal{I}_j\}) = m - 2$ as claimed.

Now, we simply split the packets and apply the optimal transmission protocol over the first halves of the packets as we did previously. This vector-linear scheme generates a SK with $m/2 - 1$ transmissions, which is an improvement of $m/2 - 1$ transmissions over the best linear scheme. Since m was arbitrary, we have shown that the gap between $S(\{\mathcal{I}_j\})$ and $S_L(\{\mathcal{I}_j\})$ cannot be bounded in general, proving Theorem 4.

Remark 3. Our proof that $S(\{\mathcal{I}_j\}) < S_L(\{\mathcal{I}_j\})$ is similar to the index coding problem, where the suboptimality of linear schemes was also shown by demonstrating a gap between the performance of linear and vector-linear coding schemes [22], [23]. For several years, it was unknown whether vector-linear coding schemes were optimal in the index coding problem. However, Blasiak et al. have since proved that even vector-linear coding is strictly suboptimal for the index coding problem [24]. We conjecture the same is true for the present setting.

IV. GENERATING MULTIPLE SECRET KEYS

Until now, we have focused exclusively on protocols that generate a single SK. However, it is also natural to consider protocols that generate τ independent secret keys. Indeed, the secrecy capacity as defined in [2] translates to the maximum number of secret keys that can possibly be generated in the combinatorial setting we consider. Thus, it is interesting to study the tradeoff between the number of secret keys that can be generated and the number of public transmissions required to do so.

To this end, we say a transmission protocol \mathbf{P} (with corresponding transmission sequence $\mathbf{T}(\underline{\mathbf{X}})$) generates τ secret keys if there exist decoding functions $\{k_1, k_2, \dots, k_n\}$ which satisfy the following three properties:

- (i) For all $j \in [n]$, and with probability 1,

$$k_j(\{X_i : i \in \mathcal{I}_j\}, \mathbf{T}(\underline{\mathbf{X}})) = k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}})).$$

- (ii) $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$ is equiprobable on \mathbb{F}^τ .

- (iii) $I(k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}})); \mathbf{T}(\underline{\mathbf{X}})) = 0$.

Note that (i)–(iii) are the same requirements for generating a single SK with one exception: we require that $k_1(\{X_i : i \in \mathcal{I}_1\}, \mathbf{T}(\underline{\mathbf{X}}))$ is uniformly over \mathbb{F}^τ . In other words, we require that each client recovers τ independent SKs, each known to all clients and private from any eavesdropper. As stated in [4, Theorem 6], Theorem 2 can be generalized as follows:

Theorem 6. For a network defined by $\{\mathcal{I}_j\}$, there exists a protocol \mathbf{P} which generates τ SKs if and only if

$$|\cup_j \mathcal{I}_j| \geq M^*(\{\mathcal{I}_j\}) + \tau. \quad (34)$$

Analogous to the definition of $S_L(\{\mathcal{I}_j\})$ in (7), let $S_L^{(\tau)}(\{\mathcal{I}_j\})$ denote the minimum number of transmissions required by a linear protocol to generate τ independent secret keys. A minor modification of our arguments for the single-SK setting yields:

Theorem 7. Let $\tau \geq 1$ be an integer. If there is a protocol \mathbf{P} for $\{\mathcal{I}_j\}$ which generates τ independent secret keys, then

$$S_L^{(\tau)}(\{\mathcal{I}_j\}) = M^*(\{\mathcal{J}_j^*\}) = |\cup_j \mathcal{J}_j^*| - \tau, \quad (35)$$

where $\{\mathcal{J}_j^*\}$ is a minimum τ -critical subfamily of $\{\mathcal{I}_j\}$.

In the single-SK setting, Lemma 4 gave a succinct interpretation of minimum 1-critical subfamilies of $\{\mathcal{I}_j\}$ as connectedness of H , the hypergraph representation of $\{\mathcal{I}_j\}$. When combined with Theorem 5,




Fig. 1: An example of a hypergraph H (left) and two induced multigraphs (center, right). Line textures are used to emphasize the relationship between the hypergraph edges and the decomposition of the multigraphs into corresponding simple connected graphs.

we find that $S_L(\{\mathcal{I}_j\})$ is in one-to-one correspondence with the size of a minimum connected dominating edge-set of H . The chief difficulty in giving a similarly succinct characterization of $S^{(\tau)}(\{\mathcal{I}_j\})$ lies in generalizing Lemma 4 appropriately for $\tau \geq 2$. In order to do so, we will need to introduce a more general notion of hypergraph connectivity.

There are many definitions of connectivity for hypergraphs. We recall two common examples here:

- *Example 1:* A hypergraph is said to be τ -edge connected if the deletion of fewer than τ edges leaves H connected.
- *Example 2:* A more stringent notion of connectivity is partition-connectivity [25]. A hypergraph H is said to be τ -partition connected if for all partitions \mathcal{P} of the vertex set, the number of hyperedges intersecting at least two parts of \mathcal{P} is at least $\tau(|\mathcal{P}| - 1)$.

We say that a multigraph is τ -partition connected if it contains τ edge-disjoint spanning trees. This definition is justified by recalling a classical result of Nash-Williams [26] and Tutte [27].

Theorem 8. *An undirected multigraph $G = (\mathcal{V}, \mathcal{E})$ contains τ edge-disjoint spanning trees iff for every partition \mathcal{P} of \mathcal{V} into disjoint sets $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_{|\mathcal{P}|}$,*

$$\sum_{e \in \mathcal{E}} (r(e; \mathcal{P}) - 1) \geq \tau(|\mathcal{P}| - 1), \quad (36)$$

where $r(e; \mathcal{P})$ is the number of parts in \mathcal{P} that the edge e intersects (i.e., its rank with respect to the partition \mathcal{P}).

Definition 2. *A multigraph $G = (\mathcal{V}, \mathcal{E}_M)$ is induced by a hypergraph $H = (\mathcal{V}, \mathcal{E})$ iff it can be decomposed into a disjoint collection of simple graphs $\{G_e\}_{e \in \mathcal{E}}$, where $G_e = (e, E_e)$ is a connected graph on the vertex set $e \in \mathcal{E}$.*

Two examples of multigraphs induced by a hypergraph are given in Figure 1.

Definition 3. *A hypergraph $H = (\mathcal{V}, \mathcal{E})$ is inherently τ -connected iff every induced multigraph contains at least τ edge-disjoint spanning trees.*




Fig. 2: A hypergraph H (left) and the induced multigraph $G_{H, \prec}$ (right) for the vertex-ordering $v_1 \prec v_2 \prec \dots \prec v_6$. Line textures are used to emphasize the relationship between the hypergraph edges and the decomposition of the multigraphs into corresponding simple connected graphs.

A pleasant generalization of Theorem 8 holds for inherently τ -connected hypergraphs.

Theorem 9. *A hypergraph $H = (\mathcal{V}, \mathcal{E})$ is inherently τ -connected iff for any partition \mathcal{P} of \mathcal{V} into disjoint sets $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_{|\mathcal{P}|}$,*

$$\sum_{e \in \mathcal{E}} (r(e; \mathcal{P}) - 1) \geq \tau(|\mathcal{P}| - 1), \quad (37)$$

where $r(e; \mathcal{P})$ is the number of parts in \mathcal{P} that the hyperedge e intersects.

Theorem 9 follows as an easy corollary of Theorem 8 and the definition of an inherently τ -connected hypergraph. However, a stronger version of Theorem 9 can be distilled from our proof of Lemma 5, which is stated shortly. Specifically, we will see that a hypergraph H is inherently τ -connected iff a relatively small subset of induced multigraphs induced by H contain τ edge-disjoint spanning trees. For a precise statement, see our remark following the proof of Lemma 5. Though not needed elsewhere in this paper, we remark that there is an analogous version of Theorem 8 for hypergraphs and partition-connectivity due to Frank, Király, and Kriesell [25].

Theorem 10. *A hypergraph $H = (\mathcal{V}, \mathcal{E})$ can be decomposed into τ subhypergraphs, each of which is 1-partition connected iff for any partition \mathcal{P} of \mathcal{V} into disjoint sets $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_{|\mathcal{P}|}$, the number of hyperedges intersecting at least two parts of \mathcal{P} is at least $\tau(|\mathcal{P}| - 1)$ (i.e., H is τ -partition connected).*

We point out that τ -partition connectivity is a more stringent condition than inherent τ -connectivity, as reflected by Theorems 9 and 10.

For a hypergraph $H = (\mathcal{V}, \mathcal{E})$, an edge set $\mathcal{E}' \subseteq \mathcal{E}$ is a minimal inherently τ -connected edge-set if the subhypergraph $H' = (\mathcal{V}, \mathcal{E}')$ is inherently τ -connected, and the removal of any edge from \mathcal{E}' results in a subhypergraph that is not inherently τ -connected. Further, define

$$\varrho_\tau(H) = \min \left\{ |\mathcal{E}'| : \mathcal{E}' \subseteq \mathcal{E} \text{ is an inherently } \tau\text{-connected edge-set} \right\}. \quad (38)$$

In other words, $\varrho_\tau(H)$ is the minimum number of edges in an inherently τ -connected subhypergraph $H' = (\mathcal{V}, \mathcal{E}')$ of $H = (\mathcal{V}, \mathcal{E})$. Note that $\varrho_\tau(H)$ is the minimum number of edges in a connected dominating edge set when $\tau = 1$, and thus its computation is **NP**-hard in general.

Lemma 5. *Let $H = (\mathcal{V}, \mathcal{E})$ be the hypergraph representation of $\{\mathcal{I}_j\}$. H is inherently τ -connected if and only if*

$$M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - \tau. \quad (39)$$

In particular, $\{\mathcal{I}_j\}$ is τ -critical if and only if \mathcal{E} is a minimal inherently τ -connected edge-set.

Before we begin the proof of Lemma 5, we take a moment to describe a special class of multigraphs that are induced by H . For a hypergraph $H = (\mathcal{V}, \mathcal{E})$, let \prec be a strict total order on \mathcal{V} . That is, if $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$, there is a permutation π on $\{1, \dots, n\}$ for which $v_{\pi(1)} \prec v_{\pi(2)} \prec \dots \prec v_{\pi(n)}$. Define the multigraph $G_{H, \prec}$ induced by H , with decomposition $\{G_e\}_{e \in \mathcal{E}}$, as follows: For each $e \in \mathcal{E}$, let G_e be a path that connects the vertices contained in e in ascending order (with respect to \prec). In other words, if $e = \{v_{i_1}, v_{i_2}, \dots, v_{i_k}\}$, where $v_{i_j} \prec v_{i_\ell}$ for $i_j < i_\ell$, then the edge-set of G_e is precisely $\{v_{i_1}, v_{i_2}\}, \{v_{i_2}, v_{i_3}\}, \dots, \{v_{i_{k-1}}, v_{i_k}\}$. An example is shown in Figure 2.

Proof of Lemma 5: Let (a_1^*, \dots, a_n^*) be an optimal solution to ILP (4). First, suppose $M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - \tau$. Then, for any partition $\mathcal{P} = \{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_k\}$ of \mathcal{V} , we have:

$$(|\cup_j \mathcal{I}_j| - \tau)(k - 1) \geq M^*(\{\mathcal{I}_j\})(k - 1) \quad (40)$$

$$= \sum_{i=1}^k \left(M^*(\{\mathcal{I}_j\}) - \sum_{j \in \mathcal{V}_i} a_j^* \right) \quad (41)$$

$$= \sum_{i=1}^k \sum_{j \in \bar{\mathcal{V}}_i} a_j^* \quad (42)$$

$$\geq \sum_{i=1}^k \left| \bigcap_{j \in \mathcal{V}_i} \bar{\mathcal{I}}_j \right| \quad (43)$$

$$= k |\cup_j \mathcal{I}_j| - \sum_{i=1}^k \left| \bigcup_{j \in \mathcal{V}_i} \mathcal{I}_j \right|, \quad (44)$$

where (43) follows by feasibility of (a_1^*, \dots, a_n^*) for ILP (4). Rearranging, we find

$$\sum_{i=1}^k \left| \bigcup_{j \in \mathcal{V}_i} \mathcal{I}_j \right| \geq |\cup_j \mathcal{I}_j| + \tau(k - 1). \quad (45)$$

Now, let G be an arbitrary multigraph induced by H with decomposition given by $\{G_e\}_{e \in \mathcal{E}}$. Note that if $e \in \mathcal{E}$ intersects $r(e; \mathcal{P})$ parts of the partition \mathcal{P} , then at least $r(e; \mathcal{P}) - 1$ edges of G_e cross the partition \mathcal{P} . Therefore,

$$\Omega(G, \mathcal{P}) \geq \sum_{e \in \mathcal{E}} (r(e; \mathcal{P}) - 1) = \left(\sum_{i=1}^k \left| \bigcup_{j \in \mathcal{V}_i} \mathcal{I}_j \right| \right) - |\cup_j \mathcal{I}_j|, \quad (46)$$

where $\Omega(G, \mathcal{P})$ denotes the number of edges in G that cross the partition \mathcal{P} . Since the partition \mathcal{P} and induced multigraph G were arbitrary, it follows from Theorem 8 and (45) that H is inherently τ -connected. Thus, we have shown:

$$M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - \tau \implies H \text{ is inherently } \tau\text{-connected.} \quad (47)$$

Next suppose H is inherently τ -connected. By optimality of (a_1^*, \dots, a_n^*) , there exists a partition $\mathcal{P}^* = \{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_k\}$ of \mathcal{V} (see [4, Appendix A], [28]) such that

$$\sum_{i=1}^k \sum_{j \in \bar{\mathcal{V}}_i} a_j^* = \sum_{i=1}^k \left| \bigcap_{j \in \mathcal{V}_i} \bar{\mathcal{I}}_j \right|. \quad (48)$$

Now, consider an arbitrary order \prec on \mathcal{V} which satisfies $u \prec v$ if $u \in \mathcal{V}_i, v \in \mathcal{V}_j$ and $i < j$. In this case, if $e \in \mathcal{E}$ intersects $r(e; \mathcal{P}^*)$ parts of the partition \mathcal{P}^* , then the path in $G_{H, \prec}$ generated by the hyperedge e (i.e., G_e) will have precisely $r(e; \mathcal{P}^*) - 1$ edges that cross \mathcal{P}^* . Since H is inherently τ -connected, we have

$$\left(\sum_{i=1}^k \left| \bigcup_{j \in \mathcal{V}_i} \mathcal{I}_j \right| \right) - |\cup_j \mathcal{I}_j| = \sum_{e \in \mathcal{E}} (r(e; \mathcal{P}^*) - 1) = \Omega(G_{H, \prec}, \mathcal{P}^*) \geq \tau(k-1) \quad (49)$$

by Theorem 8. Proceeding in a fashion similar to before, we have for \mathcal{P}^* that

$$M^*(\{\mathcal{I}_j\})(k-1) = \sum_{i=1}^k \left(M^*(\{\mathcal{I}_j\}) - \sum_{j \in \mathcal{V}_i} a_j^* \right) \quad (50)$$

$$= \sum_{i=1}^k \sum_{j \in \bar{\mathcal{V}}_i} a_j^* \quad (51)$$

$$= \sum_{i=1}^k \left| \bigcap_{j \in \mathcal{V}_i} \bar{\mathcal{I}}_j \right| \quad (52)$$

$$= k |\cup_j \mathcal{I}_j| - \sum_{i=1}^k \left| \bigcup_{j \in \mathcal{V}_i} \mathcal{I}_j \right| \quad (53)$$

$$\leq (k-1)(|\cup_j \mathcal{I}_j| - \tau), \quad (54)$$

where the final inequality follows from (49). Hence,

$$M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - \tau \iff H \text{ is inherently } \tau\text{-connected.} \quad (55)$$

We now prove the second claim. To this end, suppose $\{\mathcal{I}_j\}$ is τ -critical. Then $M^*(\{\mathcal{I}_j\}) = |\cup_j \mathcal{I}_j| - \tau$, which implies H is inherently τ -connected by (55). Consider the subhypergraph $H' = (\mathcal{V}, \mathcal{E} \setminus \{e\})$, which corresponds to the subfamily $\{\mathcal{I}_j - e\}$ of $\{\mathcal{I}_j\}$. Since $\{\mathcal{I}_j\}$ is τ -critical, we must have $M^*(\{\mathcal{I}_j - e\}) = M^*(\{\mathcal{I}_j\}) = |\cup_j \mathcal{I}_j| - \tau = |\cup_j (\mathcal{I}_j - e)| - \tau + 1$. By (55), H' cannot be inherently τ -connected, and therefore \mathcal{E} is a minimal inherently τ -connected edge-set.

On the other hand, suppose \mathcal{E} is a minimal inherently τ -connected edge-set. Then, (55) implies

$$M^*(\{\mathcal{I}_j\}) \leq |\cup_j \mathcal{I}_j| - \tau. \quad (56)$$

Since \mathcal{E} is a inherently τ -connected edge-set, for any $e \in \mathcal{E}$, $H' = (\mathcal{V}, \mathcal{E} \setminus \{e\})$ is not inherently τ -connected, and (55) implies

$$M^*(\{\mathcal{I}_j - e\}) \geq |\cup_j (\mathcal{I}_j - e)| - \tau + 1 = |\cup_j \mathcal{I}_j| - \tau. \quad (57)$$

Applying Lemma 2, we must have $M^*(\{\mathcal{I}_j\}) = M^*(\{\mathcal{I}_j - e\})$, and $|\cup_j \mathcal{I}_j| - M^*(\{\mathcal{I}_j\}) = \tau$, which implies $\{\mathcal{I}_j\}$ is τ -critical. ■

Remark 4. From the proof of Lemma 5, we observe that a hypergraph H is inherently τ -connected if and only if $G_{H, \prec}$ contains τ edge-disjoint spanning trees for every strict order \prec . Hence, this apparently weaker condition is, in fact, necessary and sufficient for any multigraph induced by H to contain τ edge-disjoint spanning trees.

In summary, we have found the following characterization of $S_L^{(\tau)}(\{\mathcal{I}_j\})$:

Theorem 11. If H is the hypergraph representation of the network defined by $\{\mathcal{I}_j\}$, then

$$S_L^{(\tau)}(\{\mathcal{I}_j\}) = \varrho_\tau(H) - \tau. \quad (58)$$

When we restrict ourselves to linear protocols, Theorem 11 elucidates a direct correspondence between the number of public transmissions required to generate τ SKs in a network and the inherent τ -connectivity of the representative hypergraph. As an illustrative example, consider the following network with 15 clients:

Example 1. Let $\mathcal{I}_1 = \{5, 7, 10, 11, 13, 14, 15\}$, and let $\{\mathcal{I}_j\}_{j=1}^{15}$ be the 14 different cyclic shifts of \mathcal{I}_1 (e.g., $\mathcal{I}_2 = \{1, 6, 8, 11, 12, 14, 15\}$, $\mathcal{I}_3 = \{1, 2, 7, 9, 12, 13, 15\}$, ...). Since the number of messages $m = 15$ is modestly small, we are able to compute $\varrho_\tau(H)$ explicitly for the hypergraph representation of the network defined by $\{\mathcal{I}_j\}$, and therefore also $S_L^{(\tau)}(\{\mathcal{I}_j\})$ by invoking Theorem 11. Below, Table I gives $S_L^{(\tau)}(\{\mathcal{I}_j\})$ for $\tau \geq 1$:

τ	1	2	3	4	5	6	≥ 7
$S_L^{(\tau)}(\{\mathcal{I}_j\})$	2	4	4	6	8	8	∞

TABLE I: $S_L^{(\tau)}(\{\mathcal{I}_j\})$ vs. τ for the network given in Example 1. Note that $S_L^{(\tau)}(\{\mathcal{I}_j\}) = \infty$ indicates that it is not possible to generate τ secret keys with any number of transmissions.

In another example⁴, we give a complete characterization for $S_L^{(\tau)}(\{\mathcal{I}_j\})$ when each pair of clients shares a unique message (i.e., $m = \binom{n}{2}$, and the hypergraph representation of the network defined by $\{\mathcal{I}_j\}$ is a complete (simple) graph on n vertices). This network model was called the PIN model by Nitinawarat and Narayan [29].

⁴We remark that this generalizes a very recent result due to Mukherjee and Kashyap [17].

Example 2. In the PIN model, $S_L^{(\tau)}(\{\mathcal{I}_j\}) = \tau(n - 2)$, where $1 \leq \tau \leq \lfloor n/2 \rfloor$. Indeed, a simple graph is inherently τ -connected iff it contains τ edge-disjoint spanning trees by Theorem 8. Thus, a simple counting argument gives $\varrho_\tau(H) = \tau(n - 1)$. An application of Theorem 11 proves the claim.

It is an interesting combinatorial design problem to specify ideal message distributions amongst clients (subject to constraints) that allow secret-key generation with fewest transmissions. For example, how many transmissions are required to generate a SK subject to the constraint that each message is initially held by at most t clients? This general problem is beyond the scope of the present paper, and we leave it for future work.

V. CONCLUDING REMARKS

In this paper, we have completely characterized the number of public transmissions required to generate a specified number of SKs when linear transmission protocols are employed. The minimum number of transmissions required by a linear protocol to generate τ secret keys is succinctly given in terms of the inherent τ -connectivity of hypergraph naturally associated with the network. We have also shown that computing said minimum number of transmissions is NP-hard.

Moreover, we have established that there can be a gap between the number of transmissions required by a nonlinear transmission scheme and the number of transmissions required by the best linear transmission scheme, and that this gap can be arbitrarily large. The problem of characterizing the number of public transmissions required by a nonlinear scheme remains an open problem, and appears to be very challenging.

REFERENCES

- [1] T. A. Courtade and T. Halford, “Coded cooperative data exchange for a secret key,” in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, Jun. 2014.
- [2] I. Csiszár and P. Narayan, “Secrecy capacities for multiple terminals,” *IEEE Trans. on Inf. Theory*, vol. 50, no. 12, pp. 3047 – 3061, Dec. 2004.
- [3] S. El Rouayheb, A. Sprintson, and P. Sadeghi, “On coding for cooperative data exchange,” in *IEEE Information Theory Workshop (ITW)*, Jan. 2010.
- [4] T. Courtade and R. Wesel, “Coded cooperative data exchange in multihop networks,” *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 1136–1158, 2014.
- [5] A. Sprintson, P. Sadeghi, G. Booker, and S. El Rouayheb, “A randomized algorithm and performance bounds for coded cooperative data exchange,” in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, June 2010, pp. 1888 –1892.
- [6] —, “Deterministic algorithm for coded cooperative data exchange,” in *ICST QShine*, Nov. 2010.
- [7] T. Courtade and R. Wesel, “Efficient universal recovery in broadcast networks,” in *Proc. of the 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct. 2010, pp. 1542 –1549.
- [8] —, “Weighted universal recovery, practical secrecy, and an efficient algorithm for solving both,” in *Proc. of the 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct. 2011.
- [9] M. Gonen and M. Langberg, “Coded cooperative data exchange problem for general topologies,” in *Proc. of IEEE International Symposium on Information Theory (ISIT)*, July 2012, pp. 2606–2610.
- [10] N. Milosavljevic, S. Pawar, S. E. Rouayheb, M. Gastpar, and K. Ramchandran, “Data exchange problem with helpers,” *arXiv preprint arXiv:1202.1612*, 2012.
- [11] N. Milosavljevic, S. Pawar, S. El Rouayheb, M. Gastpar, and K. Ramchandran, “Optimal deterministic polynomial-time data exchange for omniscience,” *arXiv preprint:1108.6046 [cs.IT]*, Aug. 2011.

- [12] S. Tajbakhsh, P. Sadeghi, and R. Shams, "A generalized model for cost and fairness analysis in coded cooperative data exchange," in *Proc. Intl. Symp. Network Coding (NetCod)*, Beijing, July 2011, pp. 1–6.
- [13] D. Ozgul and A. Sprintson, "An algorithm for cooperative data exchange with cost criterion," in *Information Theory and Applications Workshop (ITA), 2011*, Feb. 2011, pp. 1–4.
- [14] I.-H. Hou, Y.-P. Hsu, and A. Sprintson, "Truthful and non-monetary mechanism for direct data exchange," in *Proc. of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Monticello, IL, Oct. 2013.
- [15] C. Chan, "Linear perfect secret key agreement," in *2011 IEEE Information Theory Workshop Proceedings (ITW2011)*, 2011.
- [16] H. Tyagi, "Common information and secret key capacity," *IEEE Trans. on Inf. Theory*, vol. 59, no. 9, pp. 5627–5640, 2013.
- [17] M. Mukherjee and N. Kashyap, "On the communication complexity of secret key generation in the multiterminal source model," *arXiv preprint arXiv:1401.1117*, 2014.
- [18] M. Yan and A. Sprintson, "Algorithms for weakly secure data exchange," in *Network Coding (NetCod), 2013 International Symposium on*. IEEE, 2013, pp. 1–6.
- [19] S. H. Dau, W. Song, and C. Yuen, "On the existence of mds codes over small fields with constrained generator matrices," *arXiv preprint arXiv:1401.3807*, 2014.
- [20] T. Halford, T. Courtade, and K. Chugg, "Energy-efficient, secure group key agreement for ad hoc networks," in *Communications and Network Security (CNS), 2013 IEEE Conference on*, Washington, DC, Oct 2013, pp. 181–188.
- [21] M. Effros, S. El Rouayheb, and M. Langberg, "An equivalence between network coding and index coding," *arXiv:1211.6660*, Nov. 2012.
- [22] N. Alon, E. Lubetzky, U. Stav, A. Weinstein, and A. Hassidim, "Broadcasting with side information," in *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*. IEEE, 2008, pp. 823–832.
- [23] E. Lubetzky and U. Stav, "Nonlinear index coding outperforming the linear optimum," *Information Theory, IEEE Transactions on*, vol. 55, no. 8, pp. 3544–3551, 2009.
- [24] A. Blasiak, R. Kleinberg, and E. Lubetzky, "Lexicographic products and the power of non-linear network coding," in *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*. IEEE, 2011, pp. 609–618.
- [25] A. Frank, T. Király, and M. Kriesell, "On decomposing a hypergraph into i_k k_j/i_k connected sub-hypergraphs," *Discrete Applied Mathematics*, vol. 131, no. 2, pp. 373–383, 2003.
- [26] C. S. J. Nash-Williams, "Edge-disjoint spanning trees of finite graphs," *Journal of the London Mathematical Society*, vol. 1, no. 1, pp. 445–450, 1961.
- [27] W. T. Tutte, "On the problem of decomposing a graph into n connected factors," *Journal of the London Mathematical Society*, vol. 1, no. 1, pp. 221–230, 1961.
- [28] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*. Berlin: Springer-Verlag, 2003.
- [29] S. Nitinawarat and P. Narayan, "Perfect omniscience, perfect secrecy, and steiner tree packing," *Information Theory, IEEE Transactions on*, vol. 56, no. 12, pp. 6490–6500, Dec 2010.